

Reference Sequence Construction for Relative Compression of Genomes ^{*}

Shanika Kuruppu[†] Simon J. Puglisi[‡] Justin Zobel[†]

Abstract

Relative compression, where a set of similar strings are compressed with respect to a reference string, is a very effective method of compressing DNA datasets containing multiple similar sequences. Relative compression is fast to perform and also supports rapid random access to the underlying data. The main difficulty of relative compression is in selecting an appropriate reference sequence. In this paper, we explore using the dictionary of repeats generated by COMRAD, RE-PAIR and DNA-X algorithms as reference sequences for relative compression. We show this technique allows better compression and supports random access just as well. The technique also allows more general repetitive datasets to be compressed using relative compression.

1 Introduction

Rapid advancements in the field of high-throughput sequencing have led to a large number of whole genome DNA sequencing projects. Some of these projects take advantage of the improved sequencing speeds and costs, to obtain genomes of species that are unsequenced to date; for example the Genome 10K project (www.genome10k.org). Others focus on resequencing, where individual genomes from a given species are sequenced to understand variation between individuals. Examples are the 1000 Genomes project (www.1000genomes.org) for humans and the 1001 Genomes project (www.1001genomes.org) for the plant *Arabidopsis thaliana*. The assembled sequences from these projects can range from terabytes to petabytes in size. Therefore, algorithms and data structures to efficiently store, access and search these large datasets are necessary. Some progress has already been made [3, 7, 11, 17, 12], but significant challenges remain.

DNA sequences may contain repeated substrings within a sequence, however, in a database of sequences, the most significant repeats occur *between*

^{*}This work was supported by the Australian Research Council and the NICTA Victorian Research Laboratory. NICTA is funded by the Australian Government as represented by the Department of Broadband, Communications and the Digital Economy and the Australian Research Council through the ICT Center of Excellence program. *A version of this paper is to appear in the Proceedings of SPIRE 2011.*

[†]National ICT Australia, Department of Computer Science & Software Engineering, University of Melbourne, Australia, {kuruppu,jz}@csse.unimelb.edu.au

[‡]School of Computer Science and Information Technology, Royal Melbourne Institute of Technology, Australia, simon.puglisi@rmit.edu.au

sequences, usually those of the same or similar species. To help manage large genomic databases, compression algorithms that capture and efficiently encode this repeated information are employed. Compression algorithms specific to DNA sequences have been around for some time [1, 4, 5, 6, 9, 10, 19]. However, most existing algorithms are unsuitable for compressing large datasets of multiple sequences. More recently, algorithms that compress large repetitive datasets, that also support random access and search on the compressed sequences, known as *self-indexes*, have emerged. Some of these algorithms are specific to DNA compression and support random access queries [13, 14]. Others can compress general datasets and also implement search queries on the compressed sequences [11, 17].

One of the most effective ways to compress a repetitive dataset containing multiple sequences from the same or very similar species, or sequences serving the same biological functions, is to compress each sequence with respect to a chosen reference sequence [4, 14, 17]. The need for such a compression method for DNA sequences was first realised by Grumbach and Tahi [9]. XM, a statistical algorithm that implements this feature, can also generate probabilities for the level of similarity between the reference sequence and the sequence being compressed [4]. Christley, et al. proposed a solution to store just the variations of each human genome with respect to the reference genome [7] and a similar approach is taken by Brandon, et al. [3]. Mäkinen, et al. introduce more general methods to compress highly repetitive collections which also support searching in the compressed data [17].

The RLZ method, which is used in this paper, represents each sequence as an LZ77 parsing [20] with respect to a reference sequence chosen from the dataset [14]. Recently Grabowski and Deorowicz engineered RLZ to improve runtime and compression performance [8].

Relative compression algorithms like RLZ produce good compression results because the reference sequence acts as a static “dictionary” that includes most of the repeats present in the dataset being compressed. Compression speed is fast because the sequences can be compressed in a single pass over the collection, once an index on the reference sequence has been built. The static reference also makes random access fast, and easy to support. The main drawback is the difficulty of selecting an appropriate reference sequence. Selecting a reference sequence from a dataset containing only individual genomes from the same strain of the same species is simple, as any sequence will act as a good reference sequence. However this will not be effective for datasets containing sequences from different species, or from different strains of the same species.

Grabowski and Deorowicz [8] attempt to address this issue by adjusting the composition of the reference sequence during compression. When substrings of a certain minimum length, which do not occur in the reference sequence, are encountered, they are appended to the reference sequence, so that later occurrences of those substrings can be encoded as references. Results in [8] show that such a mechanism can provide a slight improvement to compression with no effects on the compression or decompression times. However, this method over-compensates and adds more substrings to the reference sequence than necessary. We compare our results with those of Grabowski and Deorowicz in Section 4.




Figure 2: The position components of the first 100 aligning factors for each sequence in the *S. cerevisiae* dataset. Only the factors that start at positions in the range of 24000-34000 are visible. The x-axis is the position on the reference sequence and the y-axis is the sequence names for sequences in the dataset.

scalable.

Moreover, a single reference sequence still may not be representative of the repetitions present in the whole dataset. A sequence may be highly similar to a few other sequences in the dataset but quite different from others. In other words, the sequences may form clusters. This is plausible for datasets containing genomes from various *strains* of a species. To test this hypothesis, we used the factors that are generated by RLZ that form alignments to the reference sequence (LISS factors that encode the segments of DNA that are not mutations [15]). We graphed the position component of these aligning factors for the *S. cerevisiae* dataset, when sequence REF is used as the reference. If the set of aligning factors are the same across two sequences, then those two sequences align to the reference sequence in the same way, hence the two sequences are similar.

The aligning factors for each sequence in the *S. cerevisiae* dataset for the position range 24,000-34,000 of the reference sequence, are illustrated in Figure 2. The graph highlights clusters of similar sequences. Most sequences have factors that start at the same position, especially those in the top half of the graph. The latter half of the graph has clusters of sequences that have similar factor positions. As an example, YPS606 and YPS128 seem to align to the reference sequence in the same way, and so do the sequences UWOPS03_461.4 and UWOPS05_227.2.

An alternative to using multiple reference sequences is to use a single reference sequence that includes the significant repeats in the whole dataset. The substrings that are shared among the sequences within clusters can be used to create a reference sequence. Dictionary compression algorithms find the repeated substrings of the dataset being compressed and the dictionary stores

these repeats. Hence, a dictionary compression algorithm that detects global repetitions can be used to generate a dictionary whose entries can then be concatenated to construct a reference sequence. We experiment with this idea next.

3 Reference Sequence Construction

We choose three dictionary compression algorithms to generate reference sequences for the two yeast datasets; RE-PAIR [16], a well-known dictionary compression algorithm, COMRAD [13], similar to RE-PAIR but tailored for DNA compression, and DNA-X [18], a DNA-specific implementation of the algorithm by Bentley and McIlroy [2]. We first compress our test datasets with RE-PAIR, COMRAD and DNA-X, and then use the dictionary of repeats as a reference sequence for relative compression. Below we explain each algorithm briefly and the process used to generate the reference sequence from the dictionary.

RE-PAIR

The RE-PAIR algorithm [16] operates in multiple iterations. In the first iteration, a count of all the distinct pairs of symbols in the input sequence are recorded. Then the most frequent symbol pair is replaced by a new symbol, and the counts are updated to reflect the replacement. In this manner, the algorithm substitutes the symbol pair with the highest count at each iteration, until there are no symbol pairs left with a count of more than one. The new symbols generated by the algorithm are identified as ‘non-terminals’, while the symbols in the original input are identified as ‘terminals’. The algorithm outputs the input sequence with all its repeated substrings replaced by non-terminal symbols, and a dictionary of rules that map all non-terminals to the symbol pairs that they replaced. The dictionary is hierarchical, since during later iterations, rules of the form $B \leftarrow CD$ or of the form $B \leftarrow cD$ or $B \leftarrow Cd$ are generated, where upper-case symbols are non-terminals and lower-case symbols are terminals. The non-terminals C and D in turn may also represent other non-terminals and so on.

The dictionary of rules generated by RE-PAIR contains the repeated substrings of the input sequence. The right hand sides of the rules can be expanded recursively to obtain the repeated substrings, which can then be concatenated to create a reference sequence. It’s not necessary to add all of the expanded rules to the reference sequence. Some of the rules lower in the hierarchy have already been incorporated into the repeated substrings of rules higher in the hierarchy that refer to these rules, so it is redundant to add these to the reference sequence. For example, expanding rule Z in the set of rules $Z \leftarrow XY$, $X \leftarrow aA$, $Y \leftarrow CD$, would result in rules X , Y , A , C and D being expanded. Once Z is expanded, it is redundant to individually expand X , Y , A , C and D . To implement this, we use a bit vector that is the length of the total number of rules. To begin with, all the bits are set to zero. When a rule Y appears on the right hand side of another rule Z , then the bit for rule Y is set to 1 to indicate that when it is Y ’s turn to get expanded later, it can be skipped.

The non-terminals generated by RE-PAIR are identified using unique integers. The higher the non-terminal number, the later the rule was generated and the higher up in the hierarchy the rule is likely to be. So starting from the highest

numbered rule to the lowest numbered rule, rule Z is expanded if and only if Z has not been expanded by a previous rule, as indicated by the bit vector. If rule Z is expanded, then the resulting substring is appended to the reference sequence. This continues until all of the rules are considered for expansion. The concatenation of the expanded substrings forms the reference sequence.

COMRAD

Similar to RE-PAIR, COMRAD [13] is a dictionary compression algorithm that detects repeated substrings in the input, and encodes them efficiently to achieve compression. COMRAD also operates in multiple iterations, however, it is a DNA-specific disk-based algorithm designed to compress large DNA datasets. Instead of replacing pairs of frequent symbols, COMRAD replaces repeated substrings of longer lengths to reduce the number of iterations.

The first iteration of COMRAD counts distinct L length substrings and the repeated substrings from most frequent to least frequent are replaced with non-terminals and a dictionary is formed. The input sequence now consists of a combination of terminals and non-terminals. In subsequent iterations, the counts of distinct substrings that satisfy a certain set of patterns is recorded (see [13]), and again substrings from most frequent to least are replaced with non-terminals. The iterations continue until there are no substrings of the above form remaining with at least a count of F (only substrings with frequency F are eligible for replacement). The algorithm outputs the input sequence with repeated substrings replaced by non-terminals, and like RE-PAIR, a dictionary containing the non-terminals mapping to the substrings they replace. As with the RE-PAIR dictionary, we expand non-terminals and append them to create a reference sequence.

DNA-X

Unlike RE-PAIR and COMRAD, DNA-X is a single pass dictionary compression algorithm. As the input is read, the fingerprint of every B -th substring of length B is stored in a hash table. To encode the next substring, all overlapping B -mers in the so far unencoded part of the input are searched for in the hash table until there is a match. The hash table gives the positions of the earlier occurrences of the B -mer. Each of these occurrences is checked to find the longest possible match. Then the prefix until the matching substring, followed by the reference for the matching substring is encoded. Searching and encoding continues until no more symbols remain to be encoded. The longest matching substrings encoded by the algorithm are the repeated substrings we use to construct the reference sequence. We modified the implementation of DNA-X by Manzini and Rastero to only output the concatenation of the longest matching substrings detected by the algorithm. We use this output as the reference sequence.

4 Experimental Results

To test the performance of the reference construction method, we use RLZ as the relative compressor. We use three test datasets containing repetitive genomes: 39 strains of *S. cerevisiae* and 36 strains of the *S. paradoxus* species of yeast, and

33 strains of *E. coli* bacteria. We ran RE-PAIR, COMRAD and DNA-X on all three datasets. For RE-PAIR, we used the default parameters, which does not place any restrictions on the number or length of repeats that can be detected. For COMRAD, we used a starting substring length L of 16 and a threshold frequency F of 2. For DNA-X we set the substring length B to 16 to be consistent with COMRAD. The repeated substrings resulting from the dictionaries were used to generate the reference sequence as described above.

Compression results are in Table 1. The first section contains the results for compressing with RLZ using the original reference sequence. The number of megabases (including the reference sequence) and the 0-order entropy of the dataset are in the first row. The second and third row contains the compression results from using the reference sequences available in the dataset with the RLZ-std and RLZ-opt (with the full set of optimisations), respectively. The results show that RLZ-opt achieves better compression compared to RLZ-std.

The second section of Table 1 contains results for using the COMRAD generated reference sequence. The two rows contain results for using the standard implementation of RLZ (RLZ-std-C) and the optimised RLZ with look-ahead and short factor encoding enabled (RLZ-opt-C)¹, respectively. The *S. cerevisiae* and *S. paradoxus* datasets compress better using the COMRAD generated reference sequence. The biggest improvement (a factor of two) is for *E. coli*. The original reference sequence was the K12 strain from the dataset, since the species does not have a reference genome. Evidently K12 is not a sequence that represents the dataset well and the COMRAD generated reference sequence is a much better representation.

The third section of Table 1 contains the results for the RE-PAIR generated reference sequences, which are very similar to the COMRAD results. The compression results improved for all three datasets with the most significant improvements being for *E. coli*. Overall, using the RE-PAIR generated reference sequences led to slightly better compressed sizes than using the COMRAD generated reference sequences.

The DNA-X generated reference sequences are not as promising. We found DNA-X generated large reference sequences, as some of the repeats it output were redundant. For example, the reference sequences for *S. cerevisiae* are 124.46 Mbases, 127.95 Mbases and 439.27 Mbases for RE-PAIR, COMRAD and DNA-X, respectively. Filtering such duplicate repeats is difficult as there are no non-terminal numbers to identify multiple occurrences of the same repeat.

Next we show that using a reference sequence containing repeats from the whole dataset is better than using a single sequence from the dataset as a reference. As in Section 2, for all three datasets, we ran RLZ-opt multiple times, with each sequence from the dataset being used as a reference at each iteration, to select a single sequence from each dataset that achieves the best compression result when used as a reference. The best compression results achieved were 9.33 MB, 13.23 MB and 18.69 MB for *S. cerevisiae* using the reference genome, *S. paradoxus* using the Z1 strain and *E. coli* using the *Sakai* strain, respectively. Comparing these results to those in the second and third sections of Table 1 shows that even if the sequence that gives the best compressed size is chosen as

¹LISS factor encoding was not used as the reference is not a sequence from the dataset and so there is no reason to expect factor positions to be predictable. For completeness, we compressed with the LISS option on and the compression results were worse than standard RLZ.

the reference sequence for a dataset, the compression results are still worse than the results that could be achieved by using a COMRAD or RE-PAIR generated reference sequence. This confirms that a single sequence is unlikely to capture all the repeats in a dataset of similar sequences and it is worth constructing a reference sequence that captures all the significant repeats of the dataset to achieve better compression results.

Table 2 shows compression and decompression times. Obviously the compression time increases significantly when using a generated reference sequence as the reference must now be generated. Also generated references tend to be longer and so more time is needed to construct suffix and LCP arrays used to perform the RLZ parsing, and to compress the reference sequence with 7zip. This is particularly the case for DNA-X. Still, performance for all methods remains at an acceptable level: the two largest datasets, can be compressed in approximately 20 minutes. More importantly, decompression times are not affected at all.

Table 3 shows compression results for RLCSA, LZ-END, COMRAD, XM and RE-PAIR algorithms being used to compress the three test datasets. The results clearly show that using RLZ with the COMRAD or RE-PAIR generated dictionaries achieve much better compression than even the best results in Table 3.

While RE-PAIR generated reference sequences seem to compress the datasets a little better than those of COMRAD, resource requirements of the algorithms should be taken into account. Both COMRAD and RE-PAIR have comparable runtimes (RE-PAIR required a little over half the time of COMRAD, see Table 2). However, the main memory usage of RE-PAIR is much higher, with *S. cerevisiae* and *S. paradoxus* using approximately 12 Gb and 11 Gb, respectively. On the other hand, COMRAD only requires 277 Mb and 554 Mb for *S. cerevisiae* and *S. paradoxus*, respectively. DNA-X has the lowest resource usage, but a better process needs to be followed to extract the necessary repeats from the dictionary to get a better quality reference sequence.

We next experiment with data sets which do not contain a specific reference. These were a *Hemoglobin* dataset containing 15,199 DNA sequences of proteins that are associated with Hemoglobin, an *Influenza* dataset containing 78,041 sequences of various strains of the Influenza virus and a *Mitochondria* dataset containing 1,521 mitochondrial DNA sequences from various species. Reference sequences were generated for the datasets using COMRAD, RE-PAIR and DNA-X. The results are presented in Table 4.

The first section of Table 4 contains the performance of RLZ when the first sequence in the dataset is chosen to be the reference. We only used standard RLZ, since the reference sequences chosen were arbitrary so none of the RLZ optimisations will be an advantage to the compression. The compression results for RLZ are worse than on previous datasets where a specific reference is available.

The results in the second section of the table are for using COMRAD generated reference sequences. Compression clearly improves for all three datasets. The most significant improvement is for the *Influenza* dataset, followed by the *Hemoglobin* dataset. The *Mitochondria* dataset did not compress very well but compression still improves.

Compression also improved significantly for all datasets by using a RE-PAIR generated reference. The *Influenza* dataset had the most significant improvement, followed by *Hemoglobin*. The *Mitochondria* dataset still does not com-

Dataset	<i>S. cerevisiae</i>		<i>S. paradoxus</i>		<i>E. coli</i>	
	Size (Mbytes)	Ent. (bpb)	Size (Mbytes)	Ent. (bpb)	Size (Mbytes)	Ent. (bpb)
Original	485.87	2.18	429.27	2.12	164.90	2.00
RLZ-std	17.89	0.29	23.38	0.44	24.27	1.18
RLZ-opt	9.33	0.15	13.44	0.25	19.30	0.94
RLZ-std-C	8.20	0.14	9.64	0.18	8.70	0.42
RLZ-opt-C	7.99	0.13	9.08	0.17	8.07	0.39
RLZ-std-R	7.78	0.13	9.10	0.17	8.21	0.40
RLZ-opt-R	7.64	0.13	8.67	0.16	7.72	0.37
RLZ-std-D	9.80	0.16	13.38	0.25	11.06	0.54
RLZ-opt-D	9.64	0.16	13.01	0.24	10.57	0.51

Table 1: Compression results for using COMRAD, RE-PAIR and DNA-X generated reference sequences. The columns are, the identifiers for RLZ version used and algorithm used to generate the reference sequence, compressed size of the dataset in Megabytes (original dataset size in Megabases) and average number of bits used per base when compressed, respectively. The sections are for compression results of RLZ when using, COMRAD, RE-PAIR and DNA-X generated reference sequences, respectively. In the first section, **RLZ-opt** includes all the optimisations. In the last two sections, **RLZ-opt** only includes looking ahead and short factor encoding.

press well. The fourth section of the table contains the results of using the DNA-X generated reference. The results have improved compared to using the original reference sequence, but gains are less than with the other two algorithms.

According to Table 4, if there is enough repetitions in the dataset, it is feasible to generate a reference sequence using either RE-PAIR or COMRAD, or any other dictionary compression algorithm, that can be used by RLZ to compress any arbitrary repetitive dataset. There is no significant difference between using a COMRAD generated reference sequence over a RE-PAIR generated one, however current implementations of RE-PAIR are less scalable than COMRAD. Table 5 shows the compression and decompression times.

Finally, we compare the new results for *S. cerevisiae* and *S. paradoxus* to those obtained by Grabowski and Deorowicz [8]. The results they achieve without the improved reference sequence are 7.18 Mbytes and 9.62 Mbytes, and with the improved reference sequence are 6.94 Mbytes and 9.01 Mbytes for *S. cerevisiae* and *S. paradoxus*, respectively. Our best results are 7.64 Mbyte for *S. cerevisiae* and 8.67 Mbyte for *S. paradoxus*, using RE-PAIR, which are comparable. It may be possible to combine the techniques to achieve even better results.

5 Concluding Remarks

Relative compression is a powerful technique for compressing collections of related genomes, which are now becoming commonplace. In this paper we have shown that these genomic collections can contain clusters of sequences which are more highly related than others. We have also shown that impressive gains in compression can be achieved by exploiting these clusters. Our specific approach

Dataset	S. cerevisiae		S. paradoxus		E. coli	
	Comp.	Dec.	Comp.	Dec.	Comp.	Dec.
	(sec)	(sec)	(sec)	(sec)	(sec)	(sec)
RLZ-std	143	9	182	6	125	3
RLZ-opt	233	8	241	6	140	3
RLZ-std-C	1561	4	1619	4	588	2
RLZ-opt-C	1783	4	1832	3	658	2
RLZ-std-R	1170	4	1134	4	455	2
RLZ-opt-R	1482	4	1353	4	499	2
RLZ-std-D	2272	8	1787	7	618	4
RLZ-opt-D	2901	7	2492	7	843	4

Table 2: Compression and decompression times for COMRAD, RE-PAIR and DNA-X generated reference sequences. The columns are: the algorithm used and the time taken to compress and decompress measured in seconds, respectively. Compression times include the time taken to generate the reference sequences, where necessary.

Dataset	S. cerevisiae		S. paradoxus		E. coli	
	Size	Ent.	Size	Ent.	Size	Ent.
	(Mbytes)	(bpb)	(Mbytes)	(bpb)	(Mbytes)	(bpb)
Original	485.87	2.18	429.27	2.12	164.90	2.00
RLCSA	41.39	0.57	47.35	0.88	34.94	1.67
LZ-END	42.52	0.70	57.18	1.07	55.25	2.68
COMRAD	15.29	0.25	18.33	0.34	13.44	0.65
XM	74.53	1.26	13.17	0.25	8.82	0.43
RE-PAIR	8.85	0.15	11.75	0.22	11.89	0.58

Table 3: Compression results for the yeast and *E. coli* datasets using other compression algorithms. The first row is the original size for all datasets (size in megabases), the remaining rows are the compression performance of RLCSA, LZ-END, COMRAD, **XM** and RE-PAIR algorithms. The two columns per dataset show the size in Mbytes and the 0-order entropy (in bits per base).

has been to detect repetitions across the dataset and build an artificial “reference sequence”, relative to which the sequence is subsequently compressed. This method retains the principle advantage of relative compression: fast random access. The drawback is slower compression time, as time must now be spent finding repeats with which to generate the reference. Future work will attempt to address this problem. We also believe it may be fruitful to apply clustering algorithms to related genomes to isolate strains.

References

- [1] B. Behzadi and F. L. Fessant. DNA compression challenge revisited: A dynamic programming approach. In *Proc. 16th Annual Symposium on Combinatorial Pattern Matching (CPM'05)*, pages 190–200, 2005.
- [2] J. Bentley and D. McIlroy. Data compression using long common strings. In *Proc. Data Compression Conference (DCC'99)*, pages 287–295, 1999.

Dataset	Hemoglobin		Influenza		Mitochondria	
	Size (Mbytes)	Ent. (bpb)	Size (Mbytes)	Ent. (bpb)	Size (Mbytes)	Ent. (bpb)
Original	7.38	2.07	112.64	1.97	25.26	1.95
RLZ-std	3.81	4.13	43.65	3.10	9.31	2.95
RLZ-std-C	1.31	1.42	3.31	0.23	6.55	2.07
RLZ-opt-C	1.17	1.27	3.00	0.21	6.05	1.92
RLZ-std-R	1.32	1.43	3.00	0.21	6.69	2.12
RLZ-opt-R	1.19	1.28	2.82	0.20	6.20	1.96
RLZ-std-D	1.42	1.54	3.68	0.26	7.13	2.26
RLZ-opt-D	1.27	1.38	3.49	0.25	6.59	2.09

Table 4: Compression results for COMRAD, RE-PAIR and DNA-X generated reference sequences for compressing repetitive datasets that do not have an explicit reference sequence. The columns are, respectively, algorithm used, compressed size in Megabytes (original dataset sizes in Megabases) and average number of bits used per base. The sections are results for using RLZ with, the first sequence from each dataset as a reference sequence, and COMRAD, RLZ and DNA-X generated reference sequences, respectively. RLZ is run in standard mode, **RLZ-std** and optimised mode, **RLZ-opt**.

- [3] M. Brandon, D. Wallace, and P. Baldi. Data structures and compression algorithms for genomic sequence data. *Bioinformatics*, 25(14):1731–1738, 2009.
- [4] M. D. Cao, T. Dix, L. Allison, and C. Mears. A simple statistical algorithm for biological sequence compression. In *Proc. Data Compression Conference (DCC’07)*, pages 43–52, 2007.
- [5] X. Chen, S. Kwong, and M. Li. A compression algorithm for DNA sequences and its applications in genome comparison. In *Proc. 4th Conference on Research in Computational Molecular Biology (RECOMB’00)*, pages 107–117, 2000.
- [6] X. Chen, M. Li, B. Ma, and J. Tromp. DNACompress: fast and effective DNA sequence compression. *Bioinformatics*, 18(12):1696–1698, 2002.
- [7] S. Christley, Y. Lu, C. Li, and X. Xie. Human genomes as email attachments. *Bioinformatics*, 25(2):274–275, 2009.
- [8] S. Grabowski and S. Deorowicz. Engineering relative compression of genomes. <http://arxiv.org/abs/1103.2351v1>, 2011.
- [9] S. Grumbach and F. Tahi. Compression of DNA sequences. In *Proc. Data Compression Conference (DCC’93)*, pages 340–350, 1993.
- [10] S. Grumbach and F. Tahi. A new challenge for compression algorithms: Genetic sequences. *Information Processing & Management*, 30(6):875–886, 1994.
- [11] S. Kreft and G. Navarro. LZ77-like compression with fast random access. In *Proc. Data Compression Conference (DCC’10)*, pages 239–248, 2010.

Dataset	Hemoglobin		Influenza		Mitochondria	
	Comp.	Dec.	Comp.	Dec.	Comp.	Dec.
	(sec)	(sec)	(sec)	(sec)	(sec)	(sec)
RLZ-std	5	1	67	11	11	1
RLZ-std-C	20	1	196	1	106	1
RLZ-opt-C	21	1	235	1	107	1
RLZ-std-R	11	1	189	1	64	1
RLZ-opt-R	12	1	220	1	64	1
RLZ-std-D	14	1	247	3	69	1
RLZ-opt-D	17	1	364	3	72	1

Table 5: Compression and decompression times for COMRAD, RE-PAIR and DNA-X generated reference sequences are used for compressing repetitive datasets that do not have an explicit reference sequence. The columns are respectively, algorithm used, total time taken to compress and decompress, measured in seconds. All compression times include time taken to generate the reference sequence.

- [12] S. Kreft and G. Navarro. Self-indexing based on LZ77. In *Proc. 22nd Annual Symposium on Combinatorial Pattern Matching (CPM'11)*, 2011. To appear.
- [13] S. Kuruppu, B. Beresford-Smith, T. Conway, and J. Zobel. Iterative dictionary construction for compression of large dna datasets. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 2011. To appear.
- [14] S. Kuruppu, S. J. Puglisi, and J. Zobel. Relative Lempel-Ziv compression of genomes for large-scale storage and retrieval. In *Proc. 17th Symposium on String Processing and Information Retrieval (SPIRE'10)*, pages 201–206, 2010.
- [15] S. Kuruppu, S. J. Puglisi, and J. Zobel. Optimized relative Lempel-Ziv compression of genomes. In *Proc. 34th Australasian Computer Science Conference (ACSC'11)*, pages 91–98, 2011.
- [16] N. J. Larsson and A. Moffat. Offline dictionary-based compression. In *Proc. Data Compression Conference (DCC'99)*, pages 296–305, 1999.
- [17] V. Mäkinen, G. Navarro, J. Sirén, and N. Välimäki. Storage and retrieval of highly repetitive sequence collections. *Journal of Computational Biology*, 17(3):281–308, 2010.
- [18] G. Manzini and M. Rastero. A simple and fast DNA compressor. *Software: Practice and Experience*, 34(14):1397–1411, 2004.
- [19] E. Rivals, J. Delahaye, M. Dauchet, and O. Delgrange. A guaranteed compression scheme for repetitive DNA sequences. In *Proc. Data Compression Conference (DCC'96)*, page 453, 1996.
- [20] J. Ziv and A. Lempel. A universal algorithm for sequential data compression. *IEEE Transactions on Information Theory*, 23(3):337–343, 1977.